

Linux Virtualization

Nesting and Management

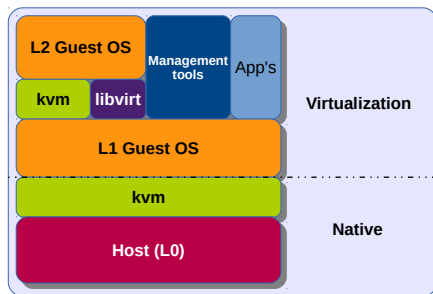
Shen Wei

Faculty of Informatics
Technische Universität München

May 14, 2013

Overview of nested virtualization

- Network bridging
- Hardware acceleration
- libvirt and virt-manager
- Guest installation



Bridge network device

Network bridging

Turning **one** network device to operate **2 or more** communication networks

- Add bridged network device
 - apt-get install bridge-utils
 - brctl addbr br0
- Configuration in **/etc/network/interfaces**

```
# allow-hotplug eth0
iface eth0 inet manual

auto br0
iface br0 inet static
    address 192.168.16.13
    netmask 255.255.255.0
    broadcast 192.168.16.255
    gateway 192.168.16.1
    bridge_ports eth0
    dns-nameservers 192.168.16.1
    dns-search tbl
```

Bridge network device

- Install bridge-utils
- Configuration in **/etc/network/interfaces**
 - Set eth0 to manual
 - Avoid (possible) network conflicts

```
# allow-hotplug eth0
iface eth0 inet manual
```

- Restart network device
 - ifdown eth0
 - ifdown br0
 - ifup br0


Hardware acceleration

- Assistance from processor: VT-x(intel), AMD-v(AMD)
 - CPU flags: vmx(intel), svm(AMD)
- Virtualization techniques for chipsets: AMD-Vi and VT-d(Intel)
- Speed up nested guest OS
 - Check CPU flags in L1 guest: vmx or svm

```
cat /proc/cpuinfo | grep flags
```
- Virtualization flags missing in guest OS
 - Consequence: **kvm.ko** not loaded
⇒ **no hardware virtualization**
 - Bug in debian squeeze?
 - No problem in debian wheezy

libvirt and virt-manager


- libvirt connects hypervisors and management applications
- Supports different hypervisors (KVM/QEMU, Xen, VMware, Virtualbox etc.)
- Unified interface for user-space VM control
Service provided by libvirtd



libvirt in virtualization (ibm.com)

libvirt and virt-manager

- libvirt connect hypervisors and management applications
- virt-manager full-featured graphical VM management tool
 - Client program of libvirtd
 - Send commands/requests
 - libvirtd handles requests and performs operations



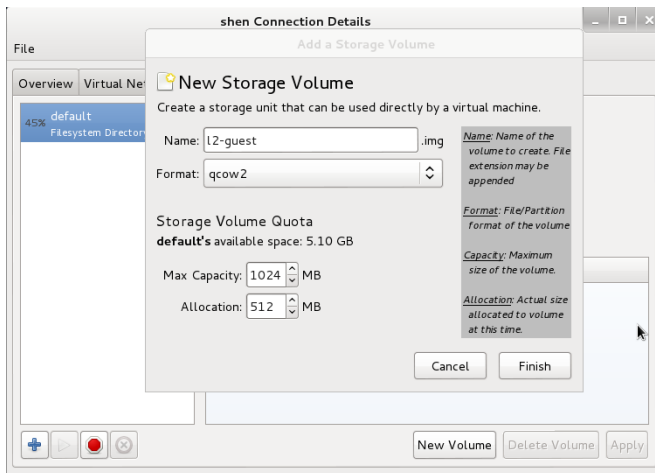
libvirt in virtualization (ibm.com)

Installation with virt-manager

- Install virt-manager: `apt-get install virt-manager`
- Create disk image
- Configuration of VM parameters
- Network settings

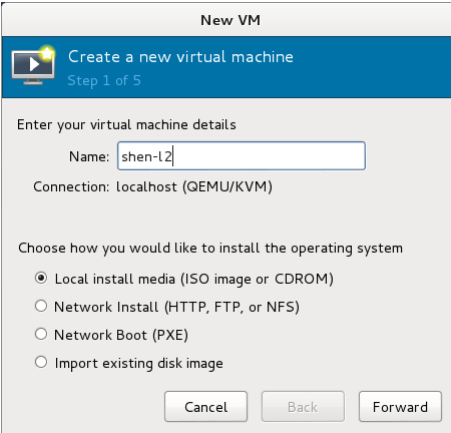
Create disk image

- virt-manager shows virtualization details from libvirtd
- Create a **Volume** in a **pool**



Register a new VM

- Choose different install medium easily
- Also possible to manage VM on remote hosts



The screenshot shows a 'New VM' dialog box. At the top, it says 'New VM' and 'Create a new virtual machine Step 1 of 5'. Below this, it asks to 'Enter your virtual machine details'. There is a text field for 'Name:' containing 'shen-l2'. Below that, it says 'Connection: localhost (QEMU/KVM)'. Then, it asks to 'Choose how you would like to install the operating system' with four radio button options: 'Local install media (ISO image or CDROM)' (selected), 'Network Install (HTTP, FTP, or NFS)', 'Network Boot (PXE)', and 'Import existing disk image'. At the bottom, there are three buttons: 'Cancel', 'Back', and 'Forward'.

New VM

Create a new virtual machine
Step 1 of 5

Enter your virtual machine details

Name: shen-l2

Connection: localhost (QEMU/KVM)

Choose how you would like to install the operating system

☒ Local install media (ISO image or CDROM)

☐ Network Install (HTTP, FTP, or NFS)


☐ Network Boot (PXE)

☐ Import existing disk image

Cancel Back Forward

RAM and processor settings

New VM

 **Create a new virtual machine**
Step 3 of 5


Choose Memory and CPU settings

Memory (RAM): MB
Up to 1002 MB available on the host

CPUs:
Up to 2 available



Assign hard drive

New VM


 **Create a new virtual machine**
Step 4 of 5

☒ Enable storage for this virtual machine

☐ Create a disk image on the computer's hard drive

  GB


5.1 Gb available in the default location

☒ Allocate entire disk now 

☒ Select managed or other existing storage

Network settings

New VM

 Create a new virtual machine
Step 5 of 5

Ready to begin installation of **shen-12**

OS: Debian Wheezy

Install: Local CDROM/ISO

Memory: 320 MB

CPUs: 1

Storage: None

☒ Customize configuration before install

▼ Advanced options

Host device eth0 (Bridge 'br0')

☒ Set a fixed MAC address

de:ad:00:00:00:43


Virt Type: kvm

Architecture: x86_64

CancelBackFinish

Display and VNC

- Use vmvga for efficient display
- Set password for VNC
- Additional options for other hardwares




Network settings in L2 guest

- Network interface during installation automatically configured (using **DHCP**)
- Manually configuration

```
allow-hotplug eth0
iface eth0 inet static
    address 192.168.16.43
    netmask 255.255.255.0
    gateway 192.168.16.1
    broadcast_ 192.168.16.255
```

VM management

- Graphical and command-line tools
- virt-manager
- virsh shell program for libvirt



Management with virsh

- Command-line user interface to libvirt
- A shell-like program:
between user and system-level component: libvirt
- Categories of virsh commands
 - domain management
 - network, interface control
 - device management
 - storage pools and volumes
 - domain monitoring
- Parameters or XML files are sent to virsh commands

A tiny example to virsh

- Define a storage pool
 - Where to store images?
 - How to access them?

```
<pool type="netfs">
  <name>netfs-pool</name>
  <source>
    <host name="192.168.16.13" />
    <dir path="/var/lib/virt/images" />
  </source>
  <target>
    <path>/var/lib/virt/images</path>
  </target>
</pool>
```

- Define a pool with virsh

```
virsh define-pool nfs-pool.xml
```


Advanced challenge:
Migration of guest VM

Migration of guest VM

- Preparation of common environment
 - Common hardware types and parameters (e.g. MAC address, PCI bus id)
 - libvirt do most tasks for us
 - Establish iSCSI storage
- Backup VM
 - Restore VM from fatal errors
- Migration to other hosts

The iSCSI standard

Internet **S**mall **C**omputer **S**ystem Interface, an IP-based networking standard for linking data storage facilities.
(Wikipedia)



<http://kensvirtualreality.wordpress.com/>

Components: target (server), initiator (client)

Preparing iSCSI storage

- iSCSI target server and initiator client
 - `apt-get install iscsitarget-dkms iscsitarget`
 - `apt-get install open-iscsi`
- Define a target in `/etc/iet/ietd.conf`
 - Restart of `ietd` required

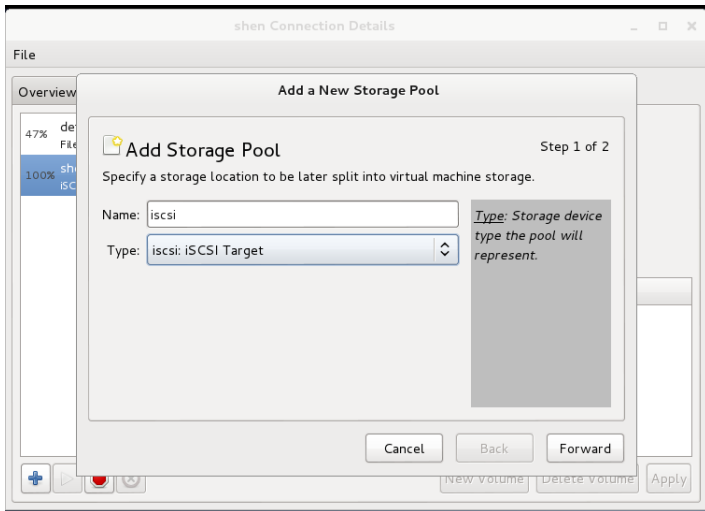
Target `iqn.2013-05.tbl.shen:debian-vm-l2`

Lun 0 Path=`/dev/disk/by-id/scsi-0QEMU_QEMU_HARDDISK_scsi1`

- Connect to iSCSI storage
 - `iscsiadm -m discovery -t st -q 192.168.16.13`
 - `iscsiadm -m node -l`

Define iSCSI pool in libvirt

- Define a pool and volume for VM



Define iSCSI pool

Add a New Storage Pool

Step 2 of 2

Add Storage Pool

Specify a storage location to be later split into virtual machine storage.

Target Path:

Host Name:

Source Path:

IQN: ☐

Source path: Path on the host that is being shared.

- Use **uuid**, **id** or **path** for more secure device mapping

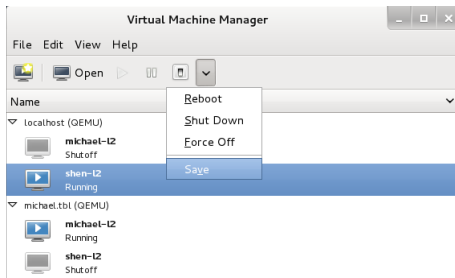
Install L2 guest VM

Install the guest VM as is described in previous slides



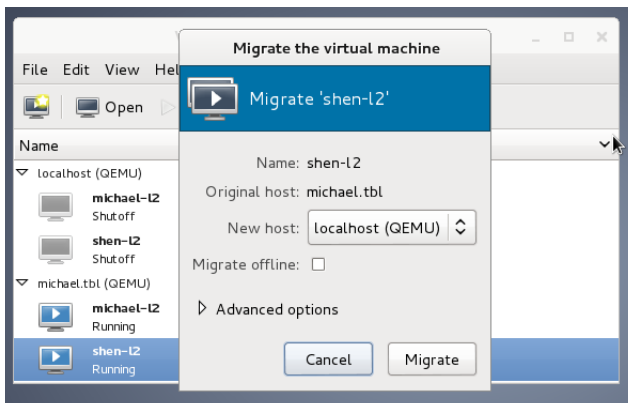
Backup guest OS

- Execute virsh with superuser account
 - virsh **save** shen-12 /home/shen/shen-12.backup
- Restore defect VM using virsh
 - virsh **restore** /home/shen/shen-12.backup
- Backup using virt-manager



Migration of VM

- Migrate to michael.tbl with virsh
virsh migrate shen-l2 \
qemu+tcp://shen@michael.tbl/system
- Migrate back to shen.tbl with virt-manager



Conclusion

- Bridged network: allow one device to serve multiple networks
- Hardware acceleration gives power to virtualization
- `libvirt` provides full featured interfaces for VM management
- Graphical and command-line tools are available for VM management
- Migration of VM's require same hardware environment

Other issues

- New version of QEMU/KVM:
 - boot option is deprecated for `-drive` parameter
 - CPU flags for virtualization support are forwarded into guest domain
- Not covered features of `libvirt`
 - Cloning VMs
 - Performance monitoring
 - ...
- Graphical install of debian system
- Set password to VNC; VNC not secure; SSH, SSL/TLS or VPN tunnelling
- `virsh` harder to learn, but powerful to play with
- **Always backup** your VM's
- Path to `ietd.conf` wrong in man page

Thank you! :)

References

- libvirt.org: Architecture of libvirt
<http://libvirt.org/intro.html>
- Wikipedia: Bridging (networking)
http://en.wikipedia.org/wiki/Bridging_%28networking%29
- wiki.debian.org: Bridging Network Connections
<http://wiki.debian.org/BridgeNetworkConnections>
- Wikipedia: X86 Virtualization
http://en.wikipedia.org/wiki/X86_virtualization
- Linux Kernel-based Virtual Machine
http://www.linux-kvm.org/page/Main_Page
- virt-manager: Virtual Machine Manager
<http://virt-manager.et.redhat.com/>

References II

- Virsh Command Reference

<http://libvirt.org/virshcmdref.html>

- XML Format for libvirt

<http://libvirt.org/format.html>

- Turning KVM

http://www.linux-kvm.org/page/Tuning_KVM

- The iSCSI Standard

<http://en.wikipedia.org/wiki/ISCSI>

- Ubuntu wiki: iSCSI (german)

<http://wiki.ubuntuusers.de/iSCSI>

- Security of Virtual Network Computing

http://en.wikipedia.org/wiki/Virtual_Network_Computing#Security